

(19)



Eur päisches Patentamt
European Patent Office
Office européen des brevets

(11)

EP 1 134 668 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
19.09.2001 Bulletin 2001/38

(51) Int. Cl. 7 : **G06F 13/40**

(21) Application number: 01300531.9

(22) Date of filing: 22.01.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Spooner, Robert John
Lower Earley, Reading Berkshire RG6 3DZ (GB)


(74) Representative:
Williams, David John et al
Page White & Farrer,
54 Doughty Street
London WC1N 2LS (GB)

(30) Priority: 15.03.2000 GB 0006291

(71) Applicant: LUCENT TECHNOLOGIES INC.
Murray Hill, New Jersey 07974-0636 (US)

(54) Data communication link between two ASICs

(57) In order to enable high speed, high bandwidth data transfer between two ASIC devices for example in a backplane, a wide parallel input data word is divided into a smaller number of words, and each smaller word is converted to serial form and then transmitted over a respective sub-link at a high clock rate relative to the system clock. At the receiving side, the clock is recovered from the serial words, and the serial words are converted back to parallel form. An alignment process is then carried out, firstly involving detecting the positions of the bits of the words and then storing the words in a buffer FIFO register. The words are clocked out of the FIFO register in synchronism under control of the system clock once it is detected that valid words are received in the FIFO registers.



Description

[0001] The present invention relates to a data communication link for high speed, high bandwidth applications.

[0002] In applications such as providing a data communication link between two Application Specific Integrated Circuits (ASICs) in a local backplane of a computing system, very high data rates may be required, e.g. an average data rate of at least 4.8 Giga bits per second (Gbps). The data link may be 64 bits wide.

[0003] Of the various possibilities for implementing such a link, it is possible to provide an interface that transfers data from the transmitting ASIC to the receiving ASIC as a single parallel word with a synchronising clock signal running at the system clock rate CK, say 78 MHz. However, for a data word of 64 bits to achieve a data transfer rate of 4.8 Gbps this would require 65 device pins, which for many applications would be either impractical or too costly to provide in the ASICs.

[0004] A synchronous interface could be used using a smaller number of pins, by multiplexing a 64 bit wide data word N times onto W bits ($= 64/N$) and by providing a synchronising clock. However with a clock signal running at 78 MHz, the bandwidth would be reduced to $W \cdot CK = BW/N$, which would give an unacceptably slow data transfer rate.

[0005] In order to achieve a bandwidth of 4.8 Gbps, the transfer rate may be multiplied N times. A synchronous interface which has a resultant Transfer Clock, $N \cdot CK$, of less than 200 MHz may be practical. Above 200 MHz, which would be necessary to achieve the desired transfer rate of 4.8 Gbps, each data bit would be valid for a maximum of 5 ns, reducing further when rise-fall times of the interconnect signals and input/output buffers are included. The task of achieving a robust design, ensuring that all W bits are aligned such that the synchronising clock can always capture valid data bytes at the receiving ASIC, is far from trivial.

Summary of the Invention

[0006] With a view to avoiding the above noted problems, the invention provides a data communication link for connection between two Application Specific Integrated Circuit (ASIC) devices that is capable of transferring wide data words at high speed. In the invention, these wide data words are considered as existing internally within each ASIC as a number of smaller sub-data words in parallel. Therefore;

$$DW = W \cdot N$$

Where;

DW = Bit width of wide data word

W = Bit width of sub-data word

N = An integer value, greater than 1

The data bandwidth across the link is given as;

$$BW = DW \cdot CK$$

Where;

BW = Bandwidth in Mega bits per second (Mbps)

CK = Transfer Clock in Mega Hertz (MHz)

For example, $W = 8$, $N = 8$, $DW = 64$. The Transfer Clock, CK, is 78 MHz giving a BW of 4992 Mbps. However the invention is not limited to these specific values.

[0007] In general terms, the invention provides a data communication link for connecting first and second devices which have a system clock, the link including transmitting means for dividing an input word into a predetermined number of smaller words, and for providing a transfer clock having a high rate relative to the system clock, and including a predetermined number of sub-links, for transmitting a respective smaller word in serial form at the transfer clock rate, and means for receiving the serial words and for converting the serial words to a parallel form, including word alignment means, responsive to the system clock, for aligning said smaller words to reconstitute the input word.

[0008] As preferred, the transmitting means includes a parallel to serial register for each smaller word, the transfer clock clocking the data from the parallel to serial register to provide high speed serial data transmission across the sub-link.

[0009] At the receive side, a serial to parallel converter is provided in each sub-link comprising a register

EP 1 134 668 A2

for receiving the incoming serial data smaller word and means responsive to the incoming data for generating a low speed clock signal therefrom, nominally equal to the frequency of the system clock, and a serial to parallel register responsive to the low speed clock for providing at an output a parallel version of the smaller word.

[0010] The low speed clock will have the same nominal frequency as the system clock, and, although aligned with the system clock to within one clock cycle, will suffer from phase jitter. In order to assemble the large number of smaller data words received into the original form of the input word, a first step which is necessary is to align the bits within each sub-link. This is effected by a bit alignment means, which comprises, on the transmit side means for sending at least two sequential initialisation words, each word having the property that no matter how many times the word is shifted left or right, there is a unique position which defines a bit alignment. Thus, by providing at the receive side a register which is at least two words wide, the position within the register of the initialisation words can be located by means of a state machine, and forwarded to the succeeding stages. The initialisation words are followed by true data words, and all the words are conveniently stored in a buffer memory, conveniently a First In First Out (FIFO) register. The words are clocked in at the recovered clock rate, but are clocked out by the system clock reference, which is approximately the same as the recovered clock, but there may be differences due to phase jitter, etc. The words are monitored as they are clocked into the FIFO register so that the words may be clocked out, when it is known all words are present, by means of the system clock as a single wide parallel word. As preferred, each FIFO register is addressed with an addressing scheme in which only one bit changes for each incremental address location. This is known as a Johnson Addressing Scheme. This permits one of the address bits to be monitored and ORed with corresponding address bits from the other FIFO registers in order to provide an alignment trigger signal when the first word is received. Since the other words will follow within one clock cycle, this trigger is provided to the output of the FIFO registers for providing a read cycle from the output of all registers.

Brief Description of the Drawings

[0011] A preferred embodiment of the invention will now be described with reference to the accompanying drawings wherein:-

Figure 1 is a schematic view of the transmit interface of a first ASIC of the data communication link of the invention;

Figure 2 is a schematic block diagram of the receive interface of a second ASIC of the data communication link of the invention;

Figure 3 is a more detailed diagram of the control mechanism for aligning received words in each sub-link of the link of Figures 1 and 2; and

Figure 4 is a schematic block diagram of a Clock Data Recovery Module (CDRM) used in both the interfaces of Figures 1 and 2.

Description of the Preferred Embodiment

[0012] Referring to Figures 1 and 2 of the drawings, a Link Interface between first and second ASICs 2, 4 includes an interface 6 in ASIC 2, which has a register 8 for breaking down the wide input data words, DW, into N (in this embodiment 8) smaller sub-words W (each 8 bits long). Each sub-word W is treated independently, using a Clock Data Recovery Module 10 (CDRM) macrocell. CDRM 10 has a multiplier 12 for multiplying the clock CK, W (8) times and respective parallel to serial (PISO) converters 14 for operating on each of N, W bit words. Each serial word is transmitted over a respective sub-link 16.

[0013] Referring to Figure 2, the receive ASIC 4 has an interface 20. The serial links 16 are coupled to another CDRM macrocell 22, in which a parallel W bit word and clock is recovered for each of the N serial links.

[0014] Figure 4 shows in more detail a CDRM 10, 22. The module 10, 22 has two primary functions. In transmit, it takes Low Speed Parallel Data (LDTX) on line 40 and creates High Speed Serial Data (HDTX) on line 42. In receive, it operates in reverse, taking High Speed Serial Data (HDRX) on line 44 and creating Low Speed Parallel Data (LDRX) on line 46. In addition, the receive operation also recovers a Low Speed Clock (LDCK) on line 48 from the serial data, that is phase aligned with the LDRX data. A Reference System Clock (REFCK) on line 49 is applied to a Phase Locked Loop 50 which multiplies the clock rate by a factor of 8 to provide a High Speed Clock (HSCK) on line 52. HSCK is applied to a parallel to serial register 54 and to a serial to parallel register 56. HSCK is also applied to a divide by 8 unit 58 and a chain of three toggles 60. The outputs of toggles 60 are detected by an edge detector

EP 1 134 668 A2

device 62 which provides an output to divider unit 58. The output of divider unit 58 comprises the Low Speed Clock (LDCK) on line 48. The operation of the circuit of Figure 4 is as follows:

[0015] For Transmit, Low Speed Data (LDTX) on line 40 will be presented to the CDRM at the rate of the Reference Clock (REFCK). The Reference Clock will be multiplied in frequency eight times by Phase Locked Loop (PLL) 50 to create High Speed Clock (HSCK) on line 52. LDTX data on line 40 will be loaded into a Parallel Serial Output (PISO) register 54 at the REFCK rate, and clocked out serially at the HSCK rate to form HDTX data on line 42.

[0016] For Receive, the High Speed Clock (HSCK) will be divided by eight at 58 to create a Low Speed Clock (LDCK) output. However, the phase of this clock must be adjusted so that its associated Low Speed Data (LDRX) is stable at the time of the active edge of LDCK. This is done by edge detection and phase adjustment circuit 60, 62 which monitors the High Speed Data (HDRX) on line 44. HDRX is also passed into a Serial Parallel Output (SIPO) register 56 to create the Low Speed Received Parallel Data (LDRX) on line 46. The output from the SIPO 56 will be enabled on the opposite edge to the active edge of its associated clock LDCK.

[0017] The number of transmit and sub-links are replicated 8 times in this example. However, there will generally only be a single PLL per CDRM macrocell.

[0018] On the receive side, the serial links are passed through CDRM macrocell 22, and a W bit word and clock will be recovered for each of the N serial links. The CDRM 22 has no knowledge of the boundary between one W bit word and the next within the serial data stream and it is therefore the first task of the Interface 20 to identify the correct bit alignment within each sub-link. Having recovered the W bit words for each sub-link, all N of the W bit words have to be aligned and synchronised to recreate the original DW width word.

[0019] The bit alignment is achieved by the transmit side sending consecutive initialisation words constructed by ASIC 2. These initialisation words (of W bits) have the property that however many times the word is shifted right or left within another word that is 2W bits wide, there is a unique position that defines the bit alignment. For example consider an initialisation word, for W = 8, of "10111000". A register 24 that is 2W words wide holds the previously received and

Previous & Current Word	Bit Alignment
10111000xxxxxxxx	0
x10111000xxxxxxxx	1
xx10111000xxxxxx	2
xxx10111000xxxxx	3
xxxx10111000xxx	4
xxxxx10111000xx	5
xxxxxx10111000x	6
xxxxxxx10111000	7

[0020] currently received words of W bits as shown in the above table. The initialisation word is sent at least twice followed by another synchronisation word (user defined) delimiter to indicate the start of transmission of true data. The position of the word is located in the register by means of a state machine (not shown) and this information is relayed to subsequent stages.

[0021] During transmission, each ASIC transmitting/receiving interface will respectively create/recreate a CRC from the true data. The CRC words are inserted at a pre-determined interval, programmed to both transmit and receive sides. After this interval the transmitted CRC should equal the recreated CRC. If not, then either bit alignment has been lost or a corruption has occurred during the transmission of the data. This provides an Integrity Check individually on each of the serial links.

[0022] Thus, as shown in Figure 2, subsequent to parallel conversion in CDRM 22, the parallel words are placed in a bit alignment register 24 in each sub-link for detecting bit alignment. This is effected by a state machine (not shown) locking onto the position of the initialisation word within the register, and passing the bit aligned word to the next stage. In the next stage, an Integrity Check is performed on the CRC word at 26.

[0023] The bit alignment and the Integrity Check are performed in each sub-link using the recovered clock generated for that serial link. There is no guarantee of any phase relationship between any of the N recovered clock (RCK[n])s, and each of the recovered clocks may be jittering (except that the recovered clocks will be within one clock cycle of one another). However, the average frequency of all recovered clocks and that of the Transfer Clock, CK, on the transmit side must be exactly the same, since the reference clock to both the transmit and receive ASICs

EP 1 134 668 A2

will be driven from the same crystal oscillator. A mechanism is therefore required to re-align the N recovered sub-words and resynchronise the wide data word back to the Transfer Clock, CK. This is done by using a short First In First Out (FIFO) 28, 6 words long, at the end of each serial link

[0024] The recovered sub-word plus a marker bit ($W + 1$ bits) is written to the FIFO 28 by its associated recovered clock on line 48. The marker bit indicates whether that data word was Transmitted Synchronisation or integrity Check Word. The very first word to be written by each of the links, will be a synchronisation word (marker bit set) and the second will be the first sub-word of true data. The first write will occur at a slightly different time for each link, but by the time the second write occurs, all will have written at least once. The addressing of the FIFOs use Johnson coding, as more clearly seen in Figure 3. An address generator 32 provides a Johnson scheme of addressing to the write read address 34 of the respective FIFO 28.

[0025] The initial value of the address is 011 and the address scheme changes as indicated in Figure 3. The most significant bit of the addresses of the sub-links are coupled by lines 36 to an OR gate 38. The output of the OR gate 38 is coupled by two metastability registers 70 to provide a trigger signal on line 72 to a state machine 74. State machine 74 provides an output on line 76 to control the reading out of the FIFO registers 28.

[0026] Thus, only a single address bit of FIFO's 28 changes per write and by ensuring that the top address bit is set on the second write, that address bit can be logically OR'd with the equivalent bit from all N links. This single bit signal, which goes high when the first word in a sub-link is received, is resynchronised via the metastability registers 70. By this time, since it is known all FIFO registers will be written to within a clock cycle of one another, all FIFO's will contain words, and the state machine 74 triggers the Word Aligner to read from all N FIFO's in parallel at the Transfer Clock rate, CK. This read should therefore occur when each of the FIFOs contain approximately four words. As the average frequency of the read and write clocks to the FIFO is the same each FIFO should always contain approximately four words. A FIFO that is at least six deep will isolate against jitter on the recovered clocks.

[0027] The very first FIFO read will all be synchronised sub-words but the second will be the recovery of the first true wide data word. The output of the FIFOs are applied to a Word Aligner register 78 which reconstitutes the original data word 80 (Figure 2). Word Alignment is checked at the same programmed interval used by the bit alignment, because at this time, and only at this time, all of the marker bits in each of the N FIFO's will be set.

[0028] The scheme outlined provides a robust high speed, high bandwidth local link by using a number of serial asynchronous links in parallel.

Claims

1. A data communication link for connecting first and second devices which have a system clock, the link including transmitting means for dividing an input word into a predetermined number of smaller words, and for providing a transfer clock having a high rate relative to the system clock, and including a predetermined number of sub-links, for transmitting a respective smaller word in serial form at the transfer clock rate, and means for receiving the serial words and for converting the serial words to a parallel form, including word alignment means, responsive to the system clock, for aligning said smaller words to reconstitute the input word.
2. A link according to claim 1, wherein the first and second devices each comprise an ASIC device, the transmitting means comprising an interface of the first ASIC device and the receiving means comprising an interface of the second ASIC device.
3. A device according to claim 1 or 2, wherein the means for transmitting comprises a parallel to serial register, and including means for generating the transfer clock coupled to the serial to parallel register for clocking serial words out of the register.
4. A communication link according to any preceding claim, including transfer clock generating means comprising a Phased Locked Loop for receiving as an input the system clock.
5. A communication link according to any preceding claim, wherein the receiving means for each sub-link comprises a serial to parallel register for receiving incoming serial words and converting them to parallel form, and means for generating a low speed clock from the incoming data with a frequency nominally equal to the system clock frequency.
6. A data communication link according to claim 5, wherein the low speed clock generating means includes edge detection means for detecting incoming data and for providing an output to a dividing means for aligning the low speed clock with recovered data and for applying the same to the serial to parallel register for clocking out

parallel words from the register.

7. A data communication link according to any preceding claim, wherein the transmitting means includes means for generating bit alignment words, and the receiving means includes a bit alignment register for storing the bit alignment words in order to locate the position of the bits in the register.
8. A data communication link according to any preceding claim, wherein the receiving means includes means for generating a cyclic redundancy code word and for transmitting the same at intervals, and the receiving means including check means for checking the cyclic redundancy code word.
9. A data communication link according to any preceding claim, including a buffer memory in each sub-link for storing a predetermined number of received words, and means for reading the buffer memories in synchronism under control of the system clock in order to reconstitute the input data word.
10. A data communication link according to claim 9, wherein the buffer memories each comprises a FIFO register.
11. A data communication link according to claim 10, wherein the FIFO registers are addressed by an addressing scheme wherein only one bit of the address changes for incremental addresses.
12. A data communication link according to claim 11, wherein a predetermined bit of the address of each FIFO are compared and employed to generate a trigger signal for actuating a state machine to cause reading of the FIFO registers.

FIG. 1 *Link Interface Transmit*




FIG. 2 *Link Interface Receive*





FIG. 3 Word Aligner Trigger

FIG. 4 Clock Data Recovery Module